NEXT-GENERATION SUPERCOMPUTING

Efficient, Affordable, and Robust

In 2002, DARPA established its High-Productivity Computing Systems (HPCS) program, with a goal of revitalizing supercomputer research and markets, and incubating a new breed of fast, efficient, easier-to-use, and affordable machines. That year, DARPA made initial grants to five key players: IBM, Cray, Hewlett-Packard, Silicon Graphics, and Sun Microsystems. The companies were given 18 months to produce concept studies for next-generation supercomputers. Though the grant amounts were small – about $3 million each – the infusion still essentially represented the end of a long drought in government investment in the field.

THE STAGE IS SET FOR DARPA INVOLVEMENT

March 11, 2002, was a turning point for supercomputing – and for the history of DARPA's role as an incubator of transformational computing research and development. On that day, researchers in Yokohama, Japan, booted up a $350 million supercomputer called the Earth Simulator, built by NEC for global climate modeling. Prior to the Earth Simulator, the fastest supercomputer had been a machine at Lawrence Livermore National Laboratory in California that could perform 7.2 trillion calculations per second, or 7.2 teraflops. The Earth Simulator surpassed 35 teraflops. But the Japanese feat was more than a scientific curiosity; it was a matter of grave concern for U.S. national security and scientific competitiveness generally. "Japanese Computer Is World's Fastest, as U.S. Falls Back," read the *New York Times* headline. In a report three years later, a National Research Council panel put matters starkly: "The Japanese Earth Simulator has served as a wake-up call, reminding us that complacency can cause us to lose not only our competitive advantage but also, and more importantly, the national competence that we need to achieve our own goals."

Of course, Japanese dominance was short-lived; in late 2002, the U.S. National Nuclear Security Administration (NNSA) gave IBM a $267 million contract to build machines faster than the Earth Simulator. IBM's Blue Gene soon surpassed the Earth Simulator, hitting 70 teraflops in 2005, and nearly doubling its own record a year later. But even this remarkable feat didn't change a structural problem that clouded the future of U.S. supercomputing competitiveness. In the early 2000s the computing industry was preoccupied with manufacture of personal computers, low-end servers, and networking gear. While clusters of commodity processors can and are used for some supercomputing applications, it was clear that over the previous 15 or 20 years, the supercomputing industry had become moribund. Dozens of companies that had specialized in high-end computing had gone out of business. U.S. government investment in that industry – and into new technology – had faltered. The National Research Council report urged key U.S. government agencies that need supercomputers – principally the Department of Energy and the Department of Defense (DoD) – to not only fund basic research, but also to help rejuvenate the industry so that it would continue to produce high-end computing tools for applications across the defense, energy, scientific, and industrial spectrum.

Dr. Charles Holland is the program manager of DARPA's High-Productivity Computing Systems program.

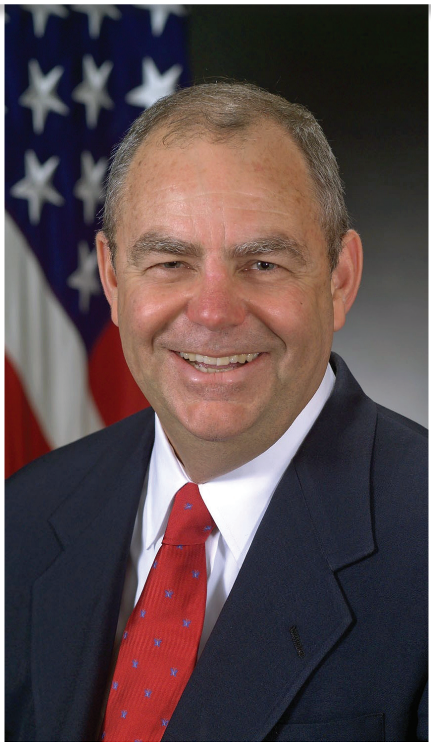

DARPA TAKES ACTION

The HPCS effort is a key part of the DoD's investment in the future of supercomputing. And the HPCS goal is to not only boost speed, but to improve computational efficiency, reduce the expense of creating and porting supercomputing applications, and make the machines more reliable. Whereas Moore's Law dictated that computing performance would double every 18 months, DARPA wanted to make sure that the supercomputer's actual value – in terms of its productivity when all factors are considered – would now double every 18 months too.

Within the DoD, the need for supercomputing is surging. The NNSA, always a key user of supercomputers, needs cutting-edge models to ensure that American nuclear weapons remain safe and operable as they age, since bombs can no longer be exploded underground for testing. New materials, weapons, and networking systems need to be proven in simulations. Warfighters need the best weather forecasting at all temporal and spatial scales. The total needs of the community are expected to surge from less than 500 teraflops in 2006 to nearly 2,500 teraflops by 2011. This community needs a computing infrastructure to sustain nearly 600 DoD projects, including climate, weather, and ocean simulations; signal and image processing; computational chemistry, biology, and materials science; environmental-quality monitoring; forces modeling and simulation; and many more.

> "We use the word 'productivity' rather than 'performance.' We are not just trying to deliver machines that just have higher performance, but that are easier to program and use."
> Dr. Charles Holland, program manager, High-Productivity Computing Systems program

Many of these DoD needs dovetail with other scientific and societal needs. The National Oceanic and Atmospheric Administration uses supercomputers to improve weather and storm surge forecasting. The Intergovernmental Panel on Climate Change uses them to model how the climate will change at various timescales and regional scales. Industries like aerospace and automotive manufacturers use models to hone designs without building physical prototypes. Biologists use them to understand how proteins fold; materials scientists use them to model nanotechnologies for next-generation materials – including ones used for computers themselves. At some level, all of these efforts enhance national security. But today, it can take months of computing time simply to model a single experiment for

Image courtesy of DARPA

Though the tightly coupled parallel systems of 1980s computing, seen here, have been far surpassed by the technological advancements of ensuing years, DARPA's High-Productivity Computing Systems program has set the bar even higher with aims to develop computing solutions that are both very productive and easier to program.

understanding the operation of nuclear fusion reactors – the perennial holy grail of power generation, which would re-create sun-like reactions to produce nearly limitless power with little waste. It is no simple task to model the behavior of materials at the atomic scale. For fusion research, as with many other kinds of research, faster, better computers can reduce research time literally by years.

To address these myriad challenges, the HPCS program will aim to push supercomputing well beyond the petaflop milestone – 1,000 teraflops, or 1,000 trillion calculations per second. But HPCS is far from a horsepower horserace, more than another Blue Gene vs. Earth Simulator. The challenge now is not just to develop computers with raw speed, but ones that are easier to use and more productive. Traditionally, supercomputers are cumbersome to program and tricky to use effectively. It can take inordinate amounts of time to design a particular application – whether a climate model or a weapons design – to work well with a given machine. One key is to optimize how you break data into many chunks, dispatch each chunk to thousands of individual processors, and then reassemble them in a timely manner. This traditionally requires time-consuming programming, customized for specific applications. And once all of this programming is finished, tuning the program to achieve good performance becomes another time-sink. If the system seems to be working more slowly than expected, it can take weeks to troubleshoot. The larger the application – and supercomputing applications are getting larger and larger – the more difficult it becomes to improve productivity, or for anyone but a supercomputer expert to figure out what might be going wrong.

"We use the word 'productivity' rather than 'performance.' We are not just trying to deliver machines that just have higher performance, but that are easier to program and use," says Charlie Holland, the program manager of the High-Productivity Computing Systems program, which is working in conjunction with the National Security Agency, the Department of Energy's Office of Science, and NNSA. "We are trying to get a way we can measure that we're getting an order of magnitude better productivity out of them." Amid an expansion of the number of areas demanding supercomputers, the machines simply need to be easier for non-experts. "We want them to be such that they are much easier to program, so people can develop codes much faster, and take codes that ran on previous machines and port them to the new machines much faster," he says. "If they are simpler to use, we can have a lot more people using these machines than used them in the past."

Among other goals, the DoD and other players would like supercomputers to be better at managing several models at once. For example, in designing a jet with "stealth" capabilities, you want both minimal radar profile and also maximum aerodynamic performance for fuel efficiency. Traditionally, modeling both fluid dynamics (the flow of air over a wing and fuselage) and modeling the electromagnetic signature (the ability of the jet to absorb or reflect radar energy) would be modeled separately. "We are trying to do designs that have a lot of parameter spaces," says Holland. "You would really like to optimize this in an integrated fashion, to optimize over many different phenomena; minimize radar signature, as well fuel efficiency. We don't have the computational power to do multiphysics design, in the time frame that would support" an aircraft design, he says. Such merged or coupled models aren't just limited to defense applications. For example, emergency planners would like to combine hurricane modeling with storm surge modeling to get a better sense of which coastal areas and buildings are most vulnerable to damage. The climate-science community would like to fully couple all ocean, land, atmosphere, and ice-sheet models to better understand how these systems interact.

As applications expand, DARPA would like to restart a robust and competitive supercomputing industry, to prime the industrial pump so that the technologies will be honed and made more affordable – and therefore more likely to be adopted. "These machines, in terms of unit of computation, are going to be a lot cheaper," adds Holland. "We didn't want to design a machine that would only work for us, and be very expensive. We really wanted a machine that would be more broadly useful. This would be cheaper and more sustainable in the long run. It is also an opportunity to enhance American competitiveness, because high-performance computing is a key enabler," he says. "What we really want them to be doing is not only make something that is useful to us at the high end," but also machines that would be commercially viable.

To help meet these and other objectives, the HPCS program set several ambitious performance goals for certain aspects of supercomputing performance. Compared with 2007's most powerful system, it

Image courtesy of DARPA

The High-Productivity Computing Systems program logo. The HPCS program has set ambitious performance goals, including a tenfold increase in execution speed over the most powerful system of 2007.

Image courtesy of DARPA

seeks an increase of more than thirtyfold in the capacity to stream data from memory to processors. It seeks a tenfold increase in execution speed. It seeks a vast increase in the capacity of the inter-node communication network (from 0.01 petabytes per second to 2.6 petabytes per second). And it seeks a demanding increase of over a thousandfold in the ability to access memory in random patterns instead of the sequential access for which systems are commonly designed. But beyond the numbers, the DARPA program seeks to address several issues to ensure that supercomputers perform well, are easily programmable, that applications can be transferred from one machine to another, that they are robust, and that all of these measures can be rendered into a clear set of productivity metrics. After two years of concept study and three years of R&D, the program is now in full development mode. IBM and Cray are the contractors for Phase III. Under the program, the two powerhouses have been tasked with meeting incremental demonstration goals, and finally demonstrating finished units by 2011.

The precise strategies of the two key players are a closely guarded secret. But while the details are under wraps, Cray and IBM are working on software technologies that make supercomputers more flexible at adapting to an application, rather than the user having to write programs specific to arcane details in the machine's hardware. Data-sharing among processors will become more automated and efficient. And on the hardware end, these more-ambitious machines will gain an improved balance of memory and network bandwidth to make efficient use of increased processing power. "IBM and Cray have different strategies," says Holland. "We're pretty early in this program. And for various reasons, a lot of the details of what their people are doing, and their designs are proprietary."

Even though the program is still under way, early payoffs are already evident. Improved Cray and IBM supercomputing technologies are already being leveraged and adapted for use by the scientific community. In 2006, the U.S. Department of Energy announced a multi-year, $200 million contract with Cray to develop the world's first petaflop-speed (1,000 teraflops) supercomputer, to be installed at the Oak Ridge National Laboratory, Oak Ridge, Tenn. By 2008, Cray is scheduled to deliver a next-generation Cray system. Code-named Baker, it is supposed to hit the 1-petaflop mark, which would make it about three times faster than any existing supercomputer in the world. The achievement is an outgrowth of DARPA's HPCS program. "They are using the technology out of this program to design and build that machine," says Holland. (Meanwhile, in 2007, Cray announced that it would provide two supercomputers to the U.S. Army Engineer Research and Development Center [ERDC] in Vicksburg, Miss. The supercomputers will give a sixfold boost to the computing capabilities of the ERDC, which supports military and civil engineering projects. The company said one of them, the Cray XT4, would hit 80 teraflops. While that's modest in comparison to the Baker machine, it would still be among the largest supercomputers in the world.)

IBM is making similarly significant advances. The National Science Foundation earlier this year announced it was awarding the University of Illinois at Urbana-Champaign's National Center for Supercomputing Applications $208 million over four-and-a-half years to install a record-shattering supercomputer based on the IBM technology. Dubbed "Blue Waters" by the university, it is expected to go online in 2011 and attain the 1-petaflop milestone. Blue Waters will be put to use for some of the most challenging modeling applications, including the impact of global warming and the evolution of galaxies. And IBM's bid was based partly on their involvement in the DARPA program.

The National Research Council, in its 2004 report, made clear the fundamental importance of supercomputing. It pointed out that "the net contributions of supercomputing, when summed over a multitude of disciplines, are no less than monumental in their impact on overall human goals. Therefore, supercomputing in some sense transcends its individual uses and can be a driver of progress in the 21st century." Thanks partly to DARPA's HPCS program, a healthy competition is already under way, new machines are being installed, and even better ones are around the corner. The program is on its way to redefining "performance" in terms of real-world "productivity" that contemplates affordability, ease of programming, and robustness. With any luck, the industry will gain new vigor to increasingly sustain supercomputing advances on its own. According to DARPA Director Dr. Anthony J. Tether, "High-productivity computing systems will give the United States a competitive advantage in the global economy. These systems will allow rapid design and test of new products and significantly shorten technology and product development timelines." And even after the HPCS program wraps up, DARPA will clearly stay involved in pushing the frontiers of supercomputing research. "We're already looking beyond this at our next-generation program," says Holland.